

AD-A012 533

BELIEF SYSTEMS AND LANGUAGE UNDERSTANDING

Bertram C. Bruce

Bolt Beranek and Newman, Incorporated

Prepared for:

Advanced Research Projects Agency

January 1975

DISTRIBUTED BY:

NTIS

National Technical Information Service
U. S. DEPARTMENT OF COMMERCE

[illegible]

Unclassified

Security Classification

AD-A 12 533

DOCUMENT CONTROL DATA - R & D

(Security classification of title, body of abstract and indexing annotation must be entered when the overall report is classified)

1. ORIGINATING ACTIVITY (Corporate author) Bolt Beranek and Newman Inc. ✓ 50 Moulton Street Cambridge, Mass. 02138		2a. REPORT SECURITY CLASSIFICATION none	
3. REPORT TITLE Belief Systems and Language Understanding		2b. GROUP	
4. DESCRIPTIVE NOTES (Type of report and inclusive dates) Technical Report		DDC RECEIVED JUL 21 1975 REGISTERED D	
5. AUTHOR(S) (First name, middle initial, last name) Bertram C. Bruce			
6. REPORT DATE January 1975	7a. TOTAL NO. OF PAGES 88	7b. NO. OF REFS 38	
8a. CONTRACT OR GRANT NO. ✓ DAHC15-71-C-0088	9a. ORIGINATOR'S REPORT NUMBER(S) ✓ BBN Report No. 2973 A.I. Report No. 21 ✓		
b. PROJECT NO. c. order no. 1697 d.	9b. OTHER REPORT NO(S) (Any other numbers that may be assigned this report)		
10. DISTRIBUTION STATEMENT Distribution of this document is unlimited. It may be released to the Clearinghouse, Department of Commerce for sale to the general public.			
11. SUPPLEMENTARY NOTES		12. SPONSORING MILITARY ACTIVITY Advanced Research Projects Agency 1400 Wilson Boulevard Arlington, Virginia 22209	

13. ABSTRACT

Understanding natural language demands the use of many types of knowledge. For example, one must know definitions of words, rules of syntax, and knowledge about the world. Some of the world knowledge has an emphasis on intention, i.e., the purpose behind actions. It is this knowledge which allows us to recognize the distinction between requests and orders, promises and threats, or lies and informings.

In order to recognize intention in behavior, whether the behavior be speech or otherwise, it is necessary to have a model of the beliefs of others. Furthermore one needs to know how actions typically fit together into larger units and how different aspects of an action are determined by intentions and beliefs.

This paper discusses some of the "belief systems knowledge" used in language understanding. It begins with a presentation of a theory of personal causation. The theory supplies the tools to account for purposeful behavior. Using primitives of the theory the social aspect of an action can be described. The social aspect is that which depends on beliefs and intentions. Patterns of behavior, called "social action paradigms" (SAP's), are then defined in terms of social actions. The SAP's provide a structure for episodes analogous to the structure a grammar provides for sentences.

PRICES SUBJECT TO CHANGE

FORM 1473
NOV 65REPLACES DD FORM 1473, 1 JAN 64, WHICH IS
OBSOLETE FOR ARMY USE.

Unclassified

Security Classification

14. KEY WORDS	LINK A		LINK B		LINK C	
	ROLE	WT	ROLE	WT	ROLE	WT
Belief Systems						
Case Structure						
Disposition Rules						
Episodes						
Intention						
Motivation Rules						
Normative Rules						
Patterns of Behavior						
Personal Causation						
Pragmatics						
Semantics						
Social Action Paradigms						
Social Actions						
Speech Acts						
Understanding Natural Language						
World Knowledge						

BELIEF SYSTEMS AND LANGUAGE UNDERSTANDING


Bertram C. Bruce

January 1975

The views and conclusions contained in this document are those of the authors and should not be interpreted as necessarily representing the official policies, either expressed or implied, of the Advanced Research Projects Agency or the U.S. Government.

This research was supported by the Advanced Research Projects Agency under ARPA Order No. 1697; Contract No. DAHC15-71-C-0088.

Distribution of this document is unlimited. It may be released to the Clearinghouse, Department of Commerce for sale to the general public.



ABSTRACT

Understanding natural language demands the use of many types of knowledge. For example, one must know definitions of words, rules of syntax, and knowledge about the world. Some of the world knowledge has an emphasis on intention, i.e., the purpose behind actions. It is this knowledge which allows us to recognize the distinction between requests and orders, promises and threats, or lies and informings.

In order to recognize intention in behavior, whether the behavior be speech or otherwise, it is necessary to have a model of the beliefs of others. Furthermore one needs to know how actions typically fit together into larger units and how different aspects of an action are determined by intentions and beliefs.

This paper discusses some of the "belief systems knowledge" used in language understanding. It begins with a presentation of a theory of personal causation. The theory supplies the tools to account for purposeful behavior. Using primitives of the theory the social aspect of an action can be described. The social aspect is that which depends on beliefs and intentions. Patterns of behavior, called "social action paradigms" (SAP's), are then defined in terms of social actions. The SAP's provide a structure for episodes analogous to the structure a grammar provides for sentences.

TABLE OF CONTENTS

	<u>page</u>
I. <u>The Role of Beliefs in Natural Language Understanding</u>	1
A. Introduction	1
B. World Knowledge	2
C. Belief Systems Knowledge	3
1. The Need for Belief Systems Knowledge	3
2. Characteristics of Belief Systems Knowledge	4
D. Previous Work	6
E. Outline of the Paper	8
1. Intentional Behavior	8
2. Social Actions	9
3. Patterns of Behavior	10
II. <u>A Theory of Personal Causation</u>	10
A. Fundamentals of the Theory	10
B. Motivation Rules	13
1. Personal Motivation Rules	13
2. Normative Rules	15
3. Disposition Rules	16
C. MUST-CHOOSE	17
D. Know and Believe	21
III. <u>Social Actions</u>	25
A. Aspects of Actions	25
B. Request	29
1. Asking ... to Verbs	29
2. Case Structure	29
3. Preconditions	30
4. Outcome Conditions	33
5. Instances and Summary Characterizations	34
6. Discussion	35
IV. <u>Patterns of Action</u>	36
A. Types of Patterns in Social Situations	36
B. Social Action Paradigms	39
C. The Social Action Paradigm *REQUEST*	44

V. <u>The Use of Belief Systems Knowledge for Language</u> <u>Understanding</u>	48
--	----

REFERENCES

ACKNOWLEDGEMENTS

APPENDIX A	A-1
APPENDIX B	B-1

I. The Role of Beliefs in Natural Language Understanding

A. Introduction

Imagine a king, for instance a King Rupert, who, on one particularly chilly morning says to his servant, Aldous,

"This place is too cool!". [1]

We have understood part of King Rupert's statement when we recognize that "this place" is the subject of the sentence, that "is too cool" is the predicate, that the predicate consists of a copula and a predicate adjective construction, and so on. Part of understanding an utterance is an at least implicit recognition of how the words in the utterance fit together. We need to distinguish nominal concepts from actions, actions from modifiers, and function words from content words before we can say that the meaning of the sentence has been discovered.

Determination of structure and word classification are clearly not enough. It is fairly safe to say that Rupert has a single thought which he is attempting to communicate by his utterance. He could mean that people in his castle are too reserved, or that the atmosphere is really hip. Most likely, though, he means that for

[1] This example was inspired by Gordon and Lakoff's (1971) Duke of Bordello.

him the temperature is too low.

B. World Knowledge

In order to recognize the sense of "cool" which Rupert intends we must apply part of our large body of knowledge about the world. We need to use such "facts" as:

SOME FACTS ABOUT THE WORLD

- A king can be expected to express his bodily wants.
- A person is uncomfortable if he is too cool.
- A king is usually conservative and does not use "hip" expressions.
- Social matters are not discussed between a king and his servant.

Many other facts are used in determining such things as how large a space is meant by "this place" and the meaning of "too" in "too cool". These facts comprise the "world knowledge" which is used in understanding unrestricted natural language. World knowledge is used to determine the context in which an utterance or written sentence occurs and to determine how that context affects the meaning of the sentence.

In this example our world knowledge tells us that Rupert means the temperature is too low in the castle, or at least in the room in which he is speaking. He also means that the low temperature is

uncomfortable for him. Only if we knew other information would we assume that he had other bad effects of the coolness in mind. For example, if Rupert were raising plants which required warm temperatures then he might be expressing his concern for the plants rather than his personal discomfort. We also know that "too cool" means something like 30-70 degrees (F) and not -200 or 3000, which might be meant in another context.

C. Belief Systems Knowledge

1. The Need for Belief Systems Knowledge

But knowing these things is still not enough. Pity the poor servant Aldous. If he knows only what has been said above about the king's statement then he may have trouble ahead. It is quite likely that Rupert is not just expressing his current state, but is in fact expressing a wish for a change. He is probably commanding Aldous to do something about the temperature, close the windows or add wood to the fire. Aldous must discover not only the literal meaning of the sentence, but also the intentions which underly it. To do that he must have additional facts such as:

ADDITIONAL FACTS ABOUT THE WORLD
(WITH AN EMPHASIS ON INTENTION)

- In this situation Rupert has authority over Aldous.
- Rupert believes that Aldous can do something about the cold.
- Rupert believes that Aldous is not necessarily planning to do something about the cold immediately.
- Rupert believes that if he orders Aldous to do something then Aldous ought to do it.

These additional facts are also part of world knowledge, but they are from a special subset which we call "belief systems knowledge". They differ from the previous set of facts only in their degree of focus on intentions, social relationships, and beliefs about other's beliefs. Like other facts, facts from a belief system can be more or less certain, complex or simple, and true or false with respect to larger bodies of knowledge. In using the term "belief systems knowledge" we are simply focusing on the kinds of world knowledge which have one or more of the following characteristics:

2. Characteristics of Belief Systems Knowledge

(1) Institutional facts (Searle (1969)) - knowledge which depends upon social conventions or commonly agreed upon rules of behavior. A baseball game exists because we have a set of agreed upon

interpretations of certain kinds of behavior. Without these interpretations we can only describe the pattern in terms of "brute facts" (Anscombe (1958)) such as "A ran from place X to place Y".

(2) Knowledge about other's knowledge - facts such as "A believes it is raining" as opposed to "it is raining". This can include higher order facts such as "A believes that B believes it is raining" or worse, "A believes that I believe it is raining". Thus part of belief systems knowledge involves what Minsky (1968) calls "models of models".

(3) Knowledge about intentions - facts about the motivations for actions. This includes procedures for assigning reasons to actions, knowledge about how plans are constructed and implemented, and knowledge of how plans of one person interact with those of another. For instance, we understand King Rupert's behavior in terms of his plan to become warm. Recognizing the plan requires knowledge about what constitutes a goal for Rupert and knowledge about the ways in which Rupert's plan affects and is affected by plans of Aldous.

D. Previous Work

The characteristics given above do not provide a sharp delineation of belief systems knowledge from other knowledge about the world. Rather they emphasize that aspect of knowledge which is concerned with human interaction, whether in discourse or other social situations. This aspect is critical for the understanding of language both because the language reflects our knowledge and because language use is itself a form of social interaction. While a precise characterization of belief systems knowledge does not exist it is worthwhile to examine what is known about belief systems and their use in language understanding. There are preliminary efforts which indicate how important knowledge of this kind can be for a system which understands language.

Within the Artificial Intelligence field the best known of these efforts are probably the works of Abelson, et al. (1965, 1969, 1973, 1975) and Colby, et al. (1969a, 1969b, 1971). Abelson's work has covered several areas of belief systems knowledge but much has focused on "ideological systems", that is, rigid, often simplistic, systems of beliefs. For example, he has studied the ideological system of the "Cold War warrior" with its beliefs about the Communist goal of world domination. The effects of new information on these systems are studied in terms of the (possibly illogical) inferences which are generated. Colby's work is similar in certain respects. He also considers the effects of new information, but emphasizes factors such as "consistency",

"foundation", and "credibility". Recently, in PARRY, a simulation of paranoid behavior, he has examined more intentionally [2] defined notions such as "fear", "mistrust", and "anger".

There has also been substantial research in linguistics, philosophy, and psychology on belief systems, work which cannot be covered in such a small space. Especially relevant are works on speech acts, presuppositions, story understanding and perception of personal causation [3]. In Artificial Intelligence the influence of a belief systems approach can be seen in works which emphasize connected discourse, context effects and pragmatics [4].

[2] Not to be confused with "intensionally".

[3] Austin (1965), Fillmore (1971), Garner (1971), Gordon and Lakoff (1971), Grice (1957), Hass (1971), Heider (1958), Hutchinson (1971), Lakoff (1971), Rumelhart (1975), Schmidt (1974), Sedlak (1974), Searle (1969), Strawson (1964), and Wittgenstein (1958).

[4] (1973), Charniak (1972), McCarthy and Hayes (1969), McDermott (1974), Minsky (1974), Norman (1972), Schank, et al, (1972a, 1972b, 1973), and Schmidt and D'Addamio (1973).

E. Outline of the Paper

This paper is not intended to be a complete survey of work on belief systems, nor is it the definitive account of what belief systems knowledge is and how it is used. Instead it is a review of two important components of a belief system, namely "social actions" and "patterns of behavior". These components rely heavily upon a recognition of intentionality in behavior and upon beliefs about one's own beliefs and the beliefs of others. Central to both is the concept of "personal causation" (Heider (1958)). A theoretical system (developed primarily by Schmidt (1974)) which accounts for intentional behavior is presented in Section II. Social actions and patterns of behavior are discussed in Sections III and IV, respectively. In Section V we consider an admittedly sketchy scenario for how these components are used in understanding a simple story. But first a brief introduction to the three components:

1. Intentional Behavior

Underlying both social actions and patterns of behavior are fundamental principles of motivated actions. These principles account for behavior which is motivated from within a person, and not accounted for by stereotyped behavior patterns or purely physical or physiological facts. The notion of a "plan" (see Schmidt (1974)) as a designed (rather than just executed) sequence of actions requires a detailed

account of how one person's beliefs are affected by the actions of others and how goal-oriented behavior takes place in a context in which other persons are also seeking goals. A sketch of this account is given in Section II. The concepts of obligation, ability, and choice discussed there are central to the definition of social actions and patterns of behavior.

2. Social Actions

A single action can be described at many levels. It may be a movement of the lips, a kiss, a sign of affection, or a kiss of death. Some of these descriptions rely on a body of knowledge about social relationships. It is only at this "social" level that we are able to describe King Rupert's utterance as a "command" rather than a "making of sound" or "simple statement". (With respect to utterances, a social level description focuses on what has been called the "speech act" (Austin (1965), Searle (1969))). In Section III and Appendix B we examine a set of social actions commonly used in understanding discourse.

3. Patterns of Behavior

Actions at any level frequently occur in patterns with other actions. These patterns of behavior (Similar to what Abelson (1974) calls "schemes" and Rumelhart (1974) calls "story schemas") are discussed in Section IV.

II. A Theory of Personal Causation*

A. Fundamentals of the Theory

Before we examine specific social action concepts and patterns of behavior we need to consider a fundamental theory of personal causation in which the various concepts can be defined. We assume that when the observer of an action in a social situation constructs his/her interpretation of the action that he/she uses a small set of recursive procedures for determining the motivation of others and a few axioms of "personal" as opposed to "natural" causation.

Certainly, behavior is explained in many ways. We often account for actions such as slipping and falling, or coughing, in physical or physiological terms. We explain other actions as being habitual or as being part of ritual behavior, such as a greeting. But not all actions can be explained in these relatively simple

*See Appendix A for notation conventions.

ways. Frequently we have to account for actions which are motivated from within a person. In those cases we need to view the action as part of a plan, i.e. a foreseeable sequence of actions leading to a goal.

Inner-motivated actions are both the most interesting and the most difficult to explain. It is necessary to appeal, not to a fixed set of causal laws, but to the beliefs of the actor, to his wants and needs, and to his beliefs about the beliefs of others. In determining the goals of the actor it is often necessary to discover his perception of the plans of others, for plans frequently depend upon the assistance or at least non-interference of others. The plan perceived by the actor may not be the same as that perceived by an observer, nor need either's perception be a correct accounting of the motivations of the third party. Furthermore, the third party's plans may depend upon his beliefs about the goals and beliefs of the original actor, the observer, or yet other persons. Clearly an account of an inner-motivated action can be quite complex. The procedures for constructing such an account must process beliefs about beliefs and plans within plans, and are hence defined recursively.

The procedures (or "motivation rules") which assign reasons to actions consider many factors, including obligation, goal-seeking and simple predisposition. They provide an account of why a person chooses to act in a given way. Together with the concept of "being able" to act, a "choose" theory is used to explain or predict

behavior. The resulting explanation is then in terms of "personal" as opposed to simple "physical" causation of events. Saying that someone "personally causes" an action implies both that he can do the action and that he is motivated to do it (where the motivation is computed by motivation rules in terms of the beliefs of the actor). This principle is called the "axiom of personal causation":

AXIOM OF PERSONAL CAUSATION
(full form)

If A CAUSES X at t

then A CAN X at t and A MUST-CHOOSE X at t

Where no confusion results we will suppress explicit indication of the times. Thus the axiom of personal causation can be written -

AXIOM OF PERSONAL CAUSATION
(abbreviated form)

If A CAUSES X then A CAN X and A MUST-CHOOSE X

The full definitions of CAUSE, CAN, and MUST-CHOOSE are given below. Informally we can read this basic axiom as "If person A does action X at time t then we can infer both that A is capable of doing X and that the weight of reasons favor A choosing to do X rather than some other action or none at all".

We say A CAN X to mean "it is possible for A to perform action X". A must have the ability and the tools, and must be in an appropriate place and position to do the act. Such "enabling conditions" are familiar from work on general problem solvers.

The CAN component of personal causation has an additional complication. It is often the case that in order to perform some action one must depend upon others to act. For example, to drive to the airport a person depends upon others driving safely, upon the gas station attendant to pump gasoline, upon the airport providing parking space, and so on. Thus CAN is necessarily defined in terms of MUST-CHOOSE, that is, MUST-CHOOSE for other actors. To account for one person's action we must consider his perception of beliefs and intentions of others.

B. Motivation Rules

The MUST-CHOOSE relation is defined in terms of relations called TRY and OUGHT. These in turn depend upon one or more "motivation rules". Motivation rules are basically procedures which relate goals of an actor to outcomes of a sequence of acts. They fall into three classes: personal, normative, and dispositional.

1. Personal Motivation Rules

Personal motivation rules account for the ways that people typically use to explain inner-motivated actions of others. These

rules are expressed in terms of the beliefs of the actor about the outcomes of his action. At least the following three rules are typically used:

(1) HEDONISM --- An action performed by A satisfies the Hedonism rule if it can be assumed that A BELIEVES that the outcomes, either immediate or eventual, of that action may include results that may be beneficial to A.

(2) EXTENDED HEDONISM --- An action performed by A satisfies the Extended Hedonism rule if it can be assumed that A BELIEVES that the outcomes, either immediate or eventual, of that action may include results that may be beneficial to someone who is in a UNIT (family, club, business partnership, etc.) or SENTIMENT (likes, loves) relationship to A, and A BELIEVES that the beneficiary WANTS the beneficial outcome.

(3) RECIPROCITY --- An action performed by A satisfies the Reciprocity rule if it can be assumed that A BELIEVES that the outcomes, either immediate or eventual, of his action constitute a "response in kind" to some action that A BELIEVES another person intentionally performed, e.g., A hit (or kissed) B

because B hit (or kissed) A.

2. Normative Rules

Normative rules (Bruce and Schmidt (1974), Schmidt (1974)) are those rules which define "conventional behavior". Whereas personal motivation rules are based upon internal motives, normative rules refer to external reasons for actions. There are specific normative rules such as:

- If someone fulfills a request you have made then you ought to thank him.

In addition there is a general normative constraint on behavior which says:

- In any ritualized situation (e.g. greeting someone, ordering food in a restaurant, or listening to a lecture) a person is expected to conform to the ritual or else be prepared to explain why not.

Generally we say that an action performed by A satisfies a normative rule if it can be assumed that A KNOWS the rule, A BELIEVES that the rule applies in this situation and if A performs the action required by that particular normative rule. Violation of a normative rule suggests inferences about the actor which cannot be made on the basis of his violation of personal motivation rules. For example, if A does not thank B for some helpful act then we may infer that A is ungrateful, rude, or perhaps, unaware. (Another interpretation

is that A does not think B's action was very helpful.)

3. Disposition Rules

The third type of motivation rule is both an extension and a mixture of the personal and normative types. Dispositional rules account for behavior which cannot be explained directly as satisfying either personal motives or normative constraints. They are invoked when an expected personal-motivation or normative rule is not followed. For example, if Fred stops to help someone fix a flat tire, we may explain his action by saying, not that Fred expected a reward or felt he should stop, but simply that Fred is helpful, where helpfulness is part of Fred's disposition. Thus "being helpful" is an extension to a large set of people of the UNIT or SENTIMENT concept in the Extended Hedonism rule. Similarly, being honest is an extension of certain normative rules. We say that an action performed by A satisfies a disposition rule if in the given situation, neither personal motivation, nor normative rules, account for the action, and the action is an instance of actions appropriate to the disposition.

C. MUST-CHOOSE

Given the various motivation rules we can define the components of MUST-CHOOSE. First there is a function on actions which gives the strength of motivation for an action. This function assumes that a motivation rule gives a "value" which is the relative strength of that motivation with respect to the given action. In a simple case the value could be 0 if the rule does not apply and 1 if it does.

Choosevalue

Choosevalue (A, X, t, S, M) --- the value of the motivation rule, M (either personal, normative, or dispositional), for the outcomes (both final and intermediate) of the sequence of acts, S, where S may be either an actual, "happened" sequence or a possible sequence; S includes A doing action X at time t. (Where no confusion results indication of times will be implicit, i.e. choosevalue (A,X,S,M).)

We can then define TRY, OUGHT, and HAS-A-REASON-TO:

TRY

(The action X has outcomes which satisfy some goal of A)

A TRY X \Leftrightarrow (forsome S) (forsome M)

[choosevalue (A, X, S, M) > 0

and (M is a personal

or personal-dispositional rule)]

OUGHT

(The action X fulfills an obligation felt by A)

A OUGHT X \Leftrightarrow (forsome S) (forsome M)
 [choosevalue (A, X, S, M) > 0
 and (M is a normative
 or normative-dispositional rule)]

HAS-A-REASON-TO

(A is motivated by either obligation or goals)

A HAS-A-REASON-TO X \Leftrightarrow (forsome S) (forsome M)
 [choosevalue (A, X, S, M) > 0]
 \Leftrightarrow A TRY X or A OUGHT X

We can now define MUST-CHOOSE as a relation on a combination of the choosevalues for a particular act. Since we are primarily interested in accounting for either (1) how an observer explains the actions in an episode, or (2) how persons use their perception of the motivations of others in forming their own plans, there is usually only one choosevalue involved. Nevertheless, the formalism allows two or more motivation rules to interact in the choice of an action.

MUST-CHOOSE

A MUST-CHOOSE X \Leftrightarrow

combine (choosevalue (A,X,S,M)) > 0

"Combine" is a function which selects an overall value for all possible sequences containing X and all possible motivation rules. In some cases "combine" may be a simple additive or maximum function, but in general may involve thresholds for reasons, interactions, and other more complex combinations of reasons.

As defined, TRY, OUGHT, and HAS-A-REASON-TO refer to acts which are done rather than not done. It is possible to choose not to act, i.e. A HAS-A-REASON-TO not-X. In that case the choosevalue must be negative. In order to account for the fact that most possible acts are not done we need to add an "axiom of laziness" which says that for any act there is some motivation not to do that act:

AXIOM OF LAZINESS

(forall X) A HAS-A-REASON-TO not-X

There is also an "axiom of negative freedom" which says that it is always possible not to do an action (occurrences like sneezing are not considered actions in this sense since choice is not involved):

AXIOM OF NEGATIVE FREEDOM

$$(\text{forall } X) A \text{ CAN not-}X$$

Together these axioms imply that

$$(\text{forall } X) A \text{ MUST-CHOOSE not-}x$$

is true, i.e. that one may always not act.

In addition to predicates which relate combinations of motives to actions it is often necessary to refer to motives which are sufficient in themselves but may not be dominating reasons in all circumstances. We say A SUFFICIENT-CHOOSE X to mean that there is a motivation rule which alone would be a sufficient reason for A to do X:

SUFFICIENT-CHOOSE

$$A \text{ SUFFICIENT-CHOOSE } X \iff (\text{forsome } S) (\text{forsome } M) \\ \text{choosevalue } (A, X, S, M) > k \text{ (where } k \text{ is a} \\ \text{threshold imposed by the axiom of laziness)}$$

Analogous to MUST-CHOOSE and SUFFICIENT-CHOOSE are relations MUST-TRY, MUST-UGHT, SUFFICIENT-TRY, and SUFFICIENT-UGHT which are restricted to certain types of motivation rules. For example, A MUST-TRY X means A MUST-CHOOSE X and M is a personal or

personal-dispositional rule.

D. Know and Believe

In addition to concepts relating to choice, a belief system requires concepts such as KNOW and BELIEVE. A sketch of these notions is given here for the sake of their use in later sections. We will consider several senses of these concepts, each defined in terms of more primitive notions.

We might begin with the Colby, et al (1969) definition of "credibility". Credibility is a function of "foundation" and "consistency" which is highest for propositions with high foundation and high consistency. "Foundation" is defined as a measure of evidence for and against a proposition. "Consistency" is a measure of the "consonance" of a proposition with other "relevant" beliefs of the individual. Credibility values range from 0 (incredible) to 100 (credible). A credibility rating of 50 means "undecided". We could define BELIEVE in terms of this credibility scale:

BELIEVE

A BELIEVE X \Leftrightarrow credibility (A, X) > 60

A possible definition for KNOW then is that KNOW is a very strong BELIEVE:

KNOW (believe-strongly)

A KNOW (believe-strongly) X \Leftrightarrow credibility (A, X) > 90

Clearly,

A KNOW (believe-strongly) X \Rightarrow A BELIEVE X

There are are other useful definitions of KNOW, however, which may not be equivalent to this one. In order to distinguish the various senses, we will use parenthetical distinguishers, e.g. KNOW (believe-strongly) for this first sense of KNOW.

A related sense of KNOW is one which separates facts which belong to the external environment from those which are inferred as belonging to the belief systems of others. For instance, a person might say "I believe that he thinks it is raining", but "I know that it is raining". Let us call this sense of KNOW, KNOW (direct). A person A might have the belief,

- If X is a belief about the beliefs of another then credibility (A, X) is necessarily less than 90

Thus no indirect belief can have a high credibility. If it were also the case that all direct beliefs had a credibility over 90 then KNOW (direct) would be equivalent to KNOW (believe-strongly).

Another useful sense of KNOW is that which distinguishes propositions believed by both the observer and the observed from those believed by just the observed person. For example, person A might say, "B knows today is St. Patrick's Day", meaning, "I believe that today is St. Patrick's Day and I believe that B believes that today is St. Patrick's. On the other hand A might say, "C believes that frogs cause warts", to mean, "I don't believe that frogs cause warts and I believe that C believes that frogs cause warts". This sense of relative KNOW is defined as follows:

KNOW (relative)

A BELIEVE (B KNOW (relative) X)

\Leftrightarrow A BELIEVE (B BELIEVE X) and A BELIEVE X

A fourth sense of KNOW which is useful is a weak sense which means that the person is aware of a proposition though he may not believe it. For example, if A tells B X we may infer that B KNOWS (is aware of) X at least for a short while following the telling. It is also true that

A BELIEVE X \Rightarrow A KNOW (is-aware-of) X

The concepts defined in this section form part of a highly interdependent theory of how persons account for the actions of others. It is closely related to language use because, in one way

or another, much of communication is concerned with such accounts. In Section III we discuss the notion of a "social action", basing definitions of specific social actions on this "Theory of Personal Causation". Several examples of social actions (especially those related to speech) are given in Appendix B. In Section IV we discuss patterns of behavior which are built out of the theory and the social actions. In Section V these concepts are applied in the analysis of a simple story.

III. Social Actions

A. Aspects of Actions

When a person utters a sentence (or writes, prints, types, etc. a sentence) he uses words to describe actions. In addition, each utterance is itself an action which can be described in words.

The description of an action can be at any of several levels, and these levels need not conflict. The idea here is analogous to that in the story of the three workers, each asked to tell what he was doing. One said "I am laying bricks"; the second said "I am building a wall"; and the third said "I am building a giant cathedral". Of course each of the workmen was right in his description of one aspect of his action. In a similar way, any act can be described at a simple physical-physiological level, or at various higher levels which take into account institutional concepts and inferred causes and effects of actions.

There are at least four aspects of actions which are important to distinguish for the design of an intelligent system. They are the physical-physiological, the propositional, the institutional, and the effectual levels. This is certainly not an exhaustive list but the implied distinctions will be sufficient to illustrate some salient characteristics of intelligent systems and of language understanding in particular.

The first aspect of an action is the physical-physiological level. For speech acts this is called the "utterance act" by Searle (1969). At this level we might describe an action as "Susie moved her arm up and down causing a paint brush to move while in contact with a chair". A speech act might have the description "Betsy uttered the sounds associated with the sentence "The Red Sox are fantastic".

The second aspect of an action is the propositional. At this level we describe actions in terms of organizing concepts. We could say "Susie is painting the chair" thus both summarizing and reinterpreting the action described above. A speech act also can be given a propositional description. Continuing our example we could say that Betsy's statement refers to the "Red Sox" and predicates "are fantastic".

The third aspect of an action is the institutional, so called because it exists by virtue of institutionalized definitions which rely on perceptions of beliefs of others. We can describe Susie's action as "helping Martha paint" if it satisfies a set of rules which constitute the definition of "help". That "help" must be defined by a set of rules about beliefs becomes clear when we consider what it is about Susie's action that makes us view it as a helping action. Certainly it is more than just the physical-physiological facts or even the propositional content of her act, for the same action could also be seen as a "harming", an "exploiting", or any of several other institutional concepts. we

have to know that Martha had a goal of painting the chair, that this goal satisfied some want or need of Martha, that Susie believed that Martha had the painting of the chair as a goal, etc.

Similarly, speech acts have institutional descriptions, or in Austin's (1965) terminology, "illocutionary force". If we believe that Betsy believes her statement, that she believes she has evidence for it, that she believes that it is not obvious to her listeners that the statement is true, that she wants her listeners to believe the statement, and perhaps other conditions, then we might describe her act as "arguing". The conditions or institutional rules which define concepts like "help" and "argue" have been called "preconditions" (Bruce and Schmidt (1974)) because they must be true at the time the concept is applied. Conditions which hold after the act has been performed are called "outcome conditions" and are used in defining the effectual aspect.

The effectual aspect is so called because it has to do with the effects or outcomes of actions. In the Betsy example, her arguing may result in her listeners becoming "convinced". In the Susie example the outcome might be that the painting is finished. For speech acts Austin calls this aspect "perlocutionary".

An action may be described at a variety of levels. As we have seen, an "uttering", a "referring and predicating", an "arguing", and a "convincing" are not different acts but different ways of conceptualizing the same act. The concept of "uttering" differs from the concept of "convincing" in that the rules for its use are

primarily physical-physiological while the rules for "convincing" have to do with the effects of an action. This is not to say that there are no physical-physiological correlates of "convincing" but only that there is a concept which summarizes a set of facts about an action; that these facts concern inferred outcomes of the action; and that the English word "convince" in its most common usage matches closely with that concept. The discussion to follow focuses on those concepts whose rules are institutional (or, as Searle says, "constitutive"). Thus we will examine the use of concepts which (unlike concepts of physical objects and actions) require a social context, a set of commonly agreed upon rules about intentionality, beliefs and social relationships, to be used and understood.

In the next section we consider the structure of a social action definition. Each concept has cases, preconditions, outcome conditions, and typical instances, or realizations in language or other behavior. Although the concepts can be defined it is important to recognize that the definitions do not imply a reduction of high level actions to primitive actions. HELP, for instance, is a social action defined in terms of beliefs and motivations, and not a complex of more primitive actions. We are able to organize a set of actions as a "helping" sequence when we infer these beliefs and motivations, but not on the basis of the action pattern itself.

B. REQUEST

1. Asking ... to Verbs

In this section we examine a social action which summarizes one person's asking another to do something. In English there are several verbs used to represent various types of asking. Austin (1965) includes these "asking...to" verbs with his "exercitives". "An exercitive is the giving of a decision in favour of or against a certain course of action, or advocacy of it." Some of the "asking...to" verbs are -- "request", "demand", "command", "beg", "order", "urge", "advise", "entreat", "warn", "plead", "direct", and "recommend". We consider the concept, REQUEST, here, and some related social action concepts in Appendix B.

REQUEST is a social action in which one person (the "agent") expresses his/her desire for another (the "recipient") to do something (the "action"). The REQUEST must, of course, be made prior to the time of the action. Unlike DEMAND and COMMAND, REQUEST does not require any commitment about moral obligations (OUGHT rules) to do the action or about explicit authority relationships.

2. Case Structure

REQUEST is defined by predicates on its various components, the persons, actions, and times. We call these components, "cases" (see Bruce (1974)). They are conceptual as opposed to grammatical relations on REQUEST. The case structure for REQUEST is represented

as follows:

REQUEST - case structure

agent: A
recipient: R
action: X
time-request: t
time-action: t'

3. Preconditions

The preconditions for REQUEST express the constraints that A intends to ask, that A wants R to do X, that A believes that R is able to do X, that A believes that R has some reason to do X, and that A believes that in the absence of the REQUEST, R will choose not to do X:

REQUEST - preconditions (full form)

- P1. A MUST-CHOOSE that
(A REQUEST R (X t) at t
- P2. A WANTS (R CAUSE X t') t
- P3. A BELIEVES (R CAN X t') t
- P4. A BELIEVES (R TRY X t') t
- P5. A BELIEVES (R MUST-CHOOSE not-X t') t
- P6. $t < t'$

Using the convention of suppressing times discussed above we can rewrite the REQUEST preconditions as follows:

REQUEST - preconditions (abbreviated form)

- P1. A MUST-CHOOSE that (A REQUEST R)
- P2. A WANTS (R CAUSE X)
- P3. A BELIEVES (R CAN X)
- P4. A BELIEVES (R TRY X)
- P5. A BELIEVES (R MUST-CHOOSE not-X)
- P6. $t < t'$

Precondition P1 is an example of a general constraint on all speech acts and on many other social actions. Basically it says that for an utterance to be a REQUESTING it must be intentional. If A is talking in his sleep, or at gunpoint, or happens to babble in a language he doesn't know then we do not describe his action as a REQUEST. This precondition alone requires the full machinery of choosing discussed in Section II. Before we recognize that a REQUEST has occurred we have to find a reason for the agent to do it.

A related precondition which might be listed for any action implying intention is one which says that the agent must BELIEVE that his intention is recognized by the recipient. An utterance is not a REQUEST if the speaker does not BELIEVE that it is perceived as a REQUEST. We can state this principle as a general condition

requiring "transparency of intention":

TRANSPARENCY CONDITION

A BELIEVE (R BELIEVE [P1])

where [P1] is the precondition on a social action which requires intentionality.

Precondition P2 says that the REQUEST occurs only when A WANTS the action to be performed. P3 and P4 state that A BELIEVES that it is possible for R to do the action and that R has some reason for doing it. But P5 says that A BELIEVES that R will not do the action without being asked. One view of a REQUEST is that it is an attempt to change R's MUST-CHOOSE value.

The transparency condition requires that A make the REQUEST obvious. Frequently, when making a REQUEST, the agent may reiterate one of the preconditions in order to satisfy this condition. Saying "I request that you..." is a way to emphasize condition P1. Any of the other preconditions can be stressed by an auxiliary statement, e.g. "And I do want you to..." or "I know you'd like to help me...".

There are other general constraints which could also be listed as preconditions on REQUEST. Searle (1969) calls these "normal input and output conditions". They include such things as that the speaker and hearer speak the same language, they are not acting in a play, and have neither aphasia nor deafness. Precondition P1 plus

the transparency condition imply most of Searle's conditions by requiring intention on the part of the speaker. The apparently simpler rule (P1) is, in fact, quite powerful because it invokes all the procedures for construction and recognition of plans.

4. Outcome Conditions

Outcome conditions on REQUEST are those facts which hold once the REQUEST has occurred (been recognized). For example, if R recognizes A's action as a REQUEST then R BELIEVES that A WANTS (R CAUSE X) and that A BELIEVES (R CAN X). Furthermore, if A BELIEVES (R BELIEVES (A REQUEST R to X)) then A BELIEVES (R BELIEVES (A WANT (R CAUSE X))) and A BELIEVES (R BELIEVES (A BELIEVES (R CAN X))). Thus A might say later, "You knew I wanted that!". Since outcomes of this sort form a potentially infinite set we will indicate that they exist by simply stating that all computable (in a weak sense determined by the situation) consequences of the preconditions and other outcome conditions are possible outcome conditions of a social action. This set of consequences will be symbolized, CONSEQUENCES (PC+OC).

The second outcome condition for REQUEST is a pointer to a normative rule. In the event that R fails to perform the action requested he is under an obligation to explain why not. The variable Z represents the reason why R does not do X.

REQUEST - Outcome Conditions

- 01. CONSEQUENCES (PC + OC)
- 02. if not-(R CAUSE X)
then R OUGHT (R EXPLAIN A Z X)

5. Instances and Summary Characterizations

Part of the definition of REQUEST is a set of pointers to typical expressions of the concept in language. We call these "instances":

REQUEST - instances

will you ..., could you ..., I would like you to ...

When an observer reports that a REQUEST has occurred he uses what we call a "summary characterization":

REQUEST - summary characterizations

ask, request

One of the interesting features of natural language is that summary characterizations can be used at the time of the action (in most cases). Thus, someone might say, "I am asking you to come with me".

This shows dramatically how belief systems knowledge interacts with both general world knowledge and specialized linguistic knowledge.

6. Discussion

There are certain verbs corresponding to social actions which cannot be used in the first person present as "ask" or "request" are used. For example "convince" sounds odd in "I am convincing you to come with me". The reason is that the concept CONVINCING has specific outcome conditions which the concept REQUEST lacks. Whenever a concept implies outcomes other than pointers to normative rules or simple consequences of the preconditions then the use of its instances in the present tense sounds strange. It is because the outcomes have obviously not occurred at the time of the utterance. This is one of the distinctions which Austin (1965) suggests between illocutionary and perlocutionary verbs.

Other "asking...to" concepts are similar to REQUEST. Most have the same case structure. Some, like THREATEN, though, have additional cases. If A THREATENS X to R there is an implied case which contains the action which A does not want R to carry out.

The preconditions and outcome conditions for the various "asking...to" concepts vary considerably, although they each contain the elements of A WANTING R to do something and A BELIEVING that R CAN do it. Definitions for several of these concepts (COMMAND, DEMAND, REQUEST, THREATEN, WARN) are given in Appendix B.

Recognition of an action at the social level facilitates understanding by providing discourse (as opposed to sentence) level constructs. Although a REQUEST may be realized by a single utterance, it is in fact a summarizing notion. Accepting such a summary commits the observer of a sequence of actions to several propositions about the beliefs and wants of the actors involved (these propositions being expressed in the preconditions and outcome conditions of REQUEST). In order to accept the propositions, the observer must attribute plans to the actors. These plans connect many actions (whether actual or foreseen) into coherent sequences. In real world situations, connections based on motivations are necessary to select useful sequences out of the many possible ones determined by physical connections alone.

IV. Patterns of Action

A. Types of Patterns in Social Situations

One way in which an intelligent system can cope with a multitude of actions performed by others is to store a set of patterns of behavior. These patterns represent summaries of many sequences of action; thus, they are efficient but occasionally faulty mechanisms for predicting or accounting for behavior.

Before discussing patterns of behavior it will be useful to reiterate the role of motivation rules in analyzing actions.

Whenever person A performs an action we can look for reasons for that action. The reason may be that A is disposed to act in that way. On the other hand one or more normative or personal motivation rules may be satisfied. As actions unfold in a sequence there is an alteration of the conditions which determine the effects of these rules. Thus previous actions play an important role in creating the environment for later choices. In the discussion to follow we consider some patterns of actions which exist because given actions in given situations facilitate (though not determine) specific successive actions.

There are a multitude of action patterns which fit the description given above. They are probably organized in some fashion which enables their efficient use in language. However it is premature to impose such an organization. Instead we will consider a few examples and some principles which should apply to the use of any such pattern.

There is one special type of behavior pattern which should be mentioned at the outset. It includes those patterns which have become ritualized, often embedded in the language. For example, the interchange -

"Hello."

"Hi! How are you?"

"Fine; and you?"

is an extremely common form of greeting. It is not difficult to imagine a grammar for such greetings which generates virtually all of the normal patterns we use. Furthermore, it is clear that the words in these patterns have lost much of their original significance. We usually do not want a graphic answer to the greeting "how are you?". Instead, the phrase is serving as an unanalyzed terminal symbol in the "greeting" grammar.

A richer example of a behavior pattern is one which might be observed in a school or other teaching and learning situation. A teacher and a student can be expected to act in a "teaching/learning" mode. This includes talking, writing on paper or a blackboard, and possibly a lot of hand waving or silences for thinking. It usually does not include physical violence, singing or love making unless one of those is the subject under study. Furthermore the content of the talking, writing, etc. is constrained. We might expect many questions and answers by both the teacher and the student. On the other hand, commands should be somewhat rare.

Included in our knowledge of the teaching/learning mode are rituals and the effects of normative and personal motivation rules. Thus this mode is a special case of very general schemas for organizing actions. A schema might be a pattern of actions without an explicit goal-orientation (like the teaching/learning pattern) or a sequence directed to a goal. In the latter case there may be only one actor involved (cf. Abelson's (1975) examples of a "plans" such

as "Communist military action against South Vietnam"). The "story schemas" of Rumelhart (1975) and the "frames" of (Minsky (1974)) are similar approaches to this issue of characterization of likely action sequences.

An important feature of any behavior pattern is that it cannot be applied in all situations. Implicit in each of the given examples are certain constraints on the applicability of the behavior pattern. The greeting pattern occurs only upon meeting someone; a waiter/customer pattern exists only in a restaurant; a boss/employee pattern exists only in the work situation; and even the teaching/learning pattern applies only when we infer certain conditions. In the following section we consider further the need for such constraints and discuss how some of them might be specified.

B. Social Action Paradigms

The term "social action paradigm" (SAP) will be used to refer to a pattern of behavior with its associated situational constraints. The situational constraints form the "header" of the social action paradigm and the pattern itself is the "body". The header checks conditions in the situation, and, in the process, determines the bindings of variables in the SAP body to entities (individuals, outcomes of acts, acts, propositions) in the situation.

Part of a SAP can be represented by a graph in which the nodes are action-concepts and the arcs represent likely actions to follow. For example a teacher and a student might engage in the interaction shown in Figure 1.

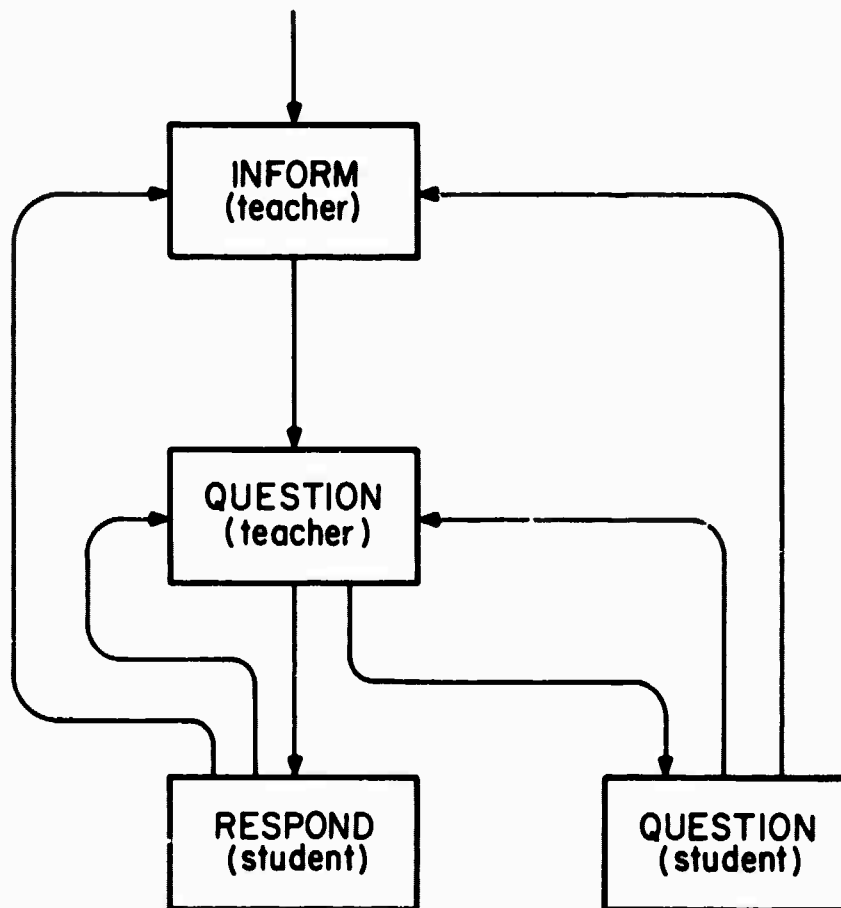


Figure 1. Rough approximation of a SAP body.

That such a network of social actions is insufficient can be seen by consideration of the relationship between the successive actions . For instance, when the teacher poses a question we expect the student to respond, but this response should be related to the question. It may not be a correct or complete answer, but to be a "response to the question" its propositional content should bear a close relationship to the propositional content of the question. If this "close relationship" is defined by a function, F , then we can give a better account of the question-response paradigm as shown in Figure 2.

Using other propositional content relating functions we can redo Figure 2 as shown in Figure 3.

Once the notion of social actions with arguments (or cases, see Bruce (1974)) is introduced the problem of specifying the situations in which the paradigm is applicable becomes explicit. That is, it is necessary to show the correspondences, or bindings, of entities in the situation to the variables in the SAP graph.

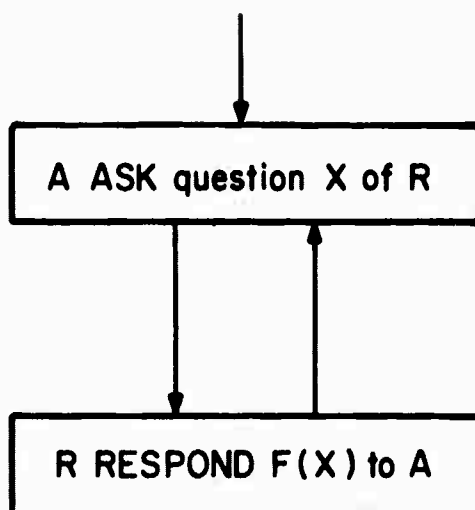


Figure 2. Instantiated social actions with a propositional content relating function, F.

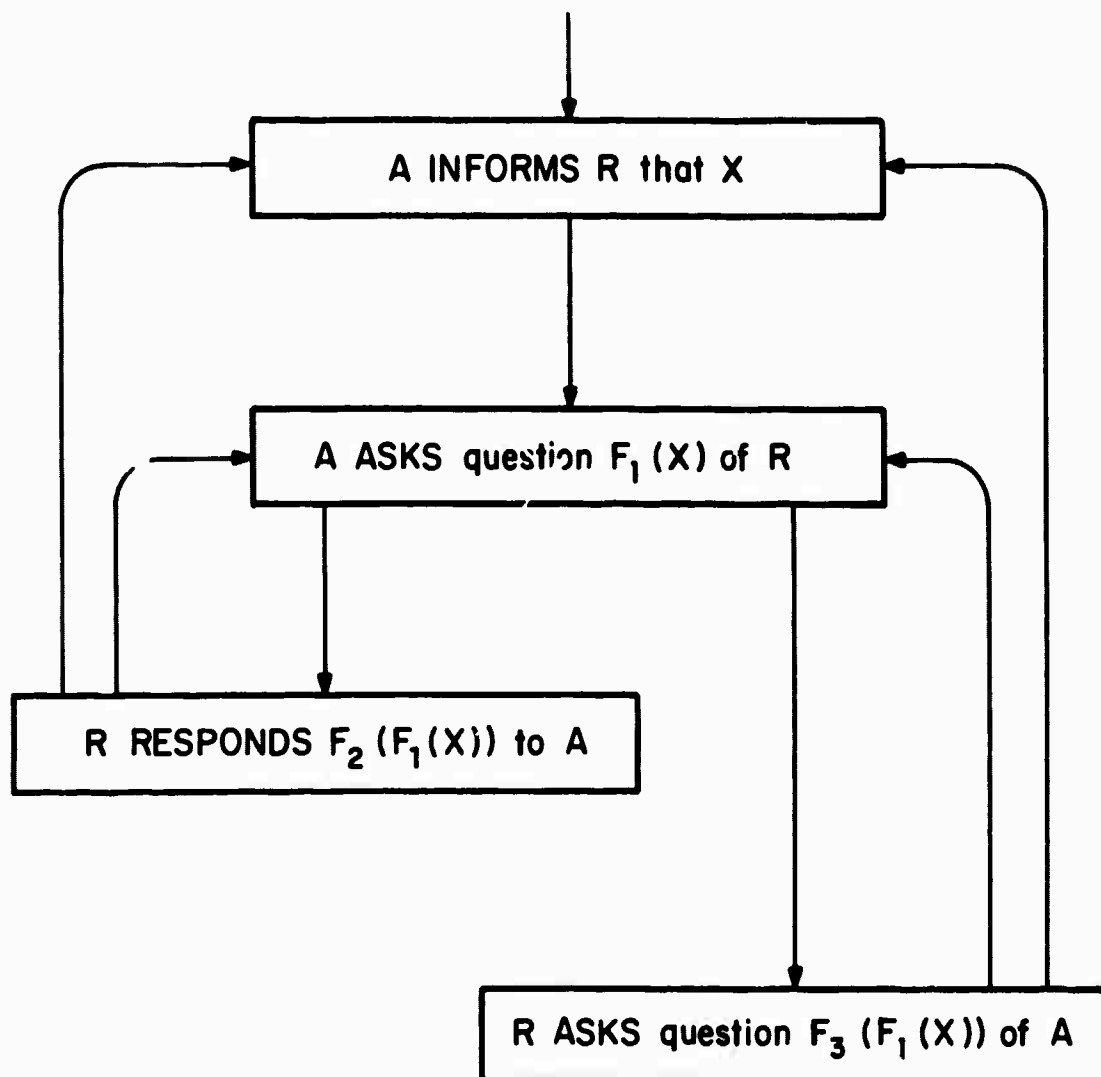


Figure 3. A SAP body with instantiated social actions.

The problem of specifying appropriate situations is, of course, inherent in a SAP whether or not the variables are explicitly identified. In the previous example it is clear that the SAP exists because there are commonly shared expectations concerning the behavior of students and teachers. Thus, whenever we observe two individuals in a student-teacher relationship, with the associated relations of dominance, helping, etc. we are able to invoke the student-teacher SAP. The SAP could certainly be applied in other situations, but most likely in situations which share certain salient aspects of the teacher-student situation. We would not, for example, expect similar behavior to occur between a ship's captain and his first mate (assuming they were not also in a teaching situation).

In the following section we consider a social action paradigm which we find associated with the concept, REQUEST.

C. The Social Action Paradigm, *REQUEST*

REQUEST is a pattern of behavior (body) with constraints (header) on its applicability. The header for *REQUEST* specifies both conditions on the situation in which *REQUEST* is to be applied and bindings of variables in the body to entities of the situation. *REQUEST* is a simple SAP in the sense that its pattern has only one starting node, that is, the pattern begins when person A REQUESTS that person R do some action X. Because of this, the header is the

same as the preconditions for REQUEST:

***REQUEST* header**

1. A MUST-CHOOSE that (A REQUEST R (R CAUSE X))
2. A WANTS (R CAUSE X)
3. A BELIEVES (R CAN X)
4. A BELIEVES (R TRY X)
5. A BELIEVES (R MUST-CHOOSE not-X)
6. time-request < time-X

If there were other starting nodes in the ***REQUEST*** body then the preconditions of the corresponding social actions would also have to be satisfied, and hence included in the ***REQUEST*** header. Other SAP's have additional constraints in their headers. For example, the teaching/learning SAP discussed above might have the constraint -

A is-teaching R

The body for ***REQUEST*** is represented by a graph in which the nodes are instantiated social actions and the arcs show which actions are likely to follow from a given action. (Figure 4.)

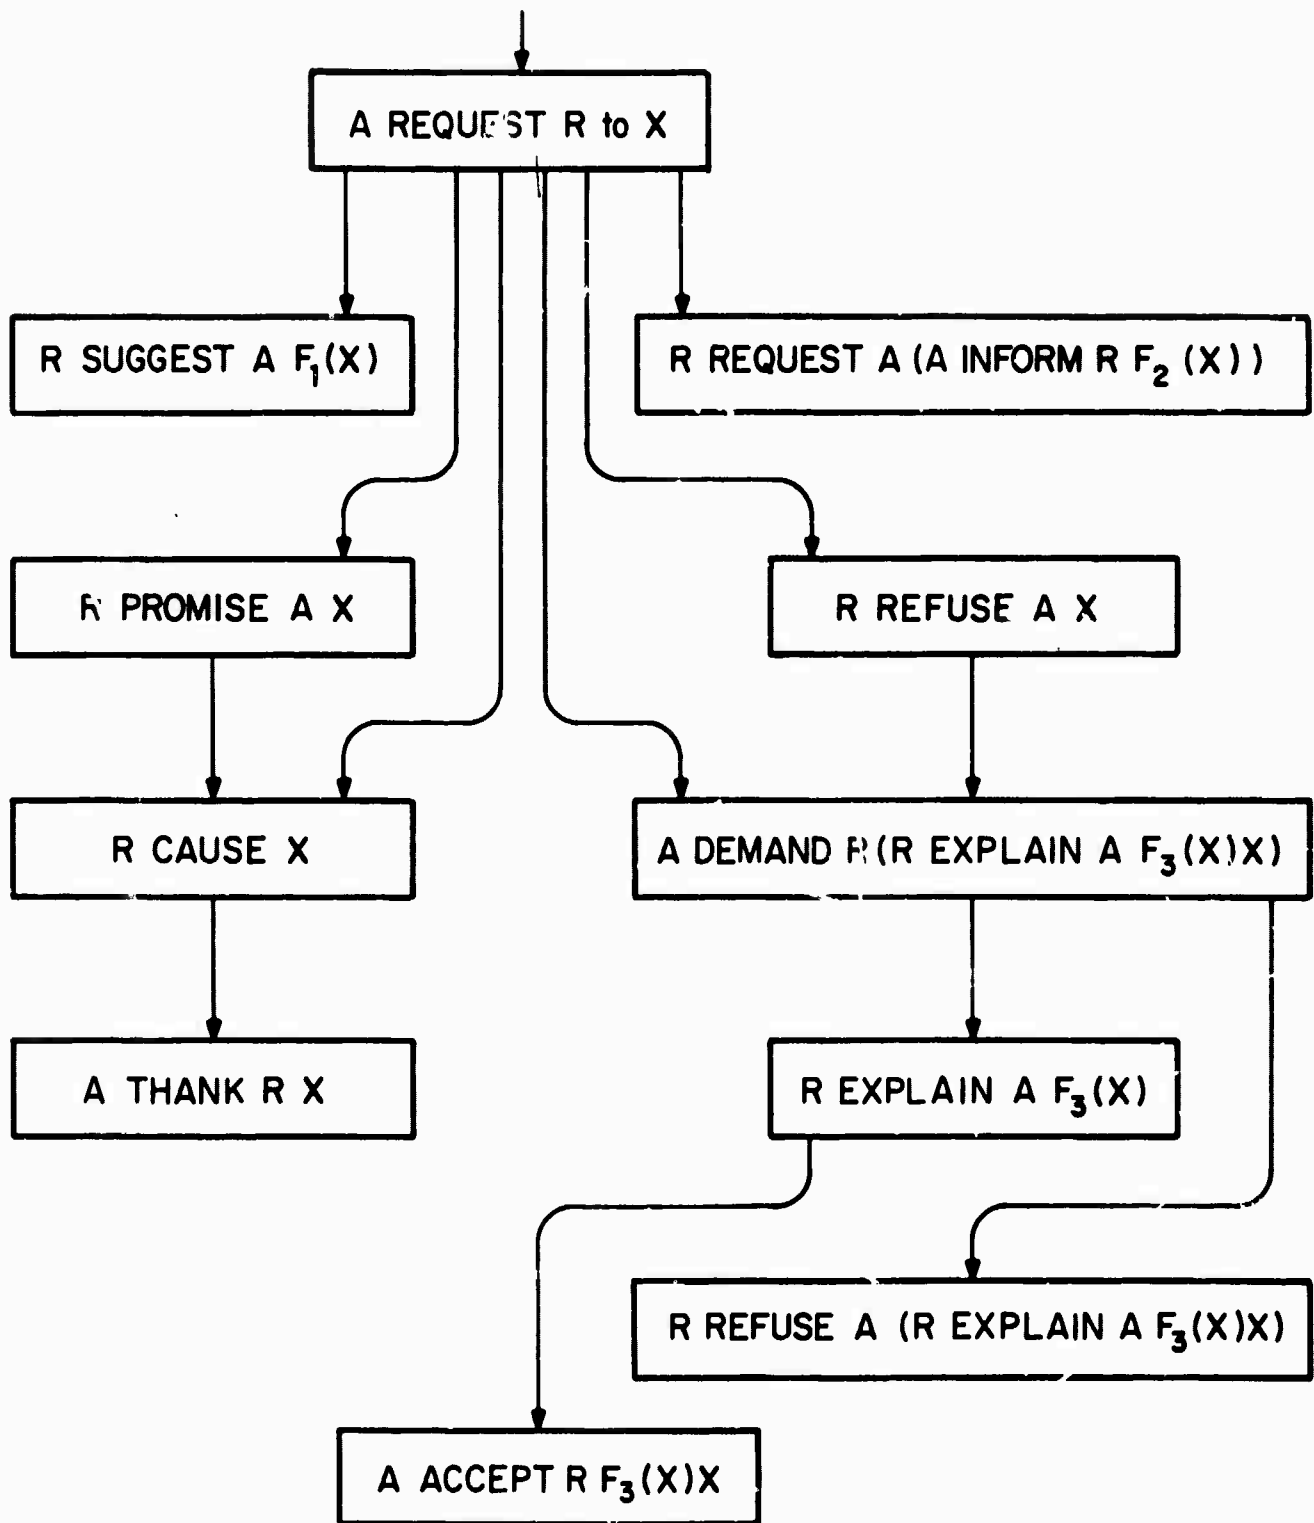


Figure 4. The social action paradigm, *REQUEST*.

In the SAP the functions F_1 , F_2 and F_3 relate the propositional content of one action to another. $F_1(X)$ is an action which can be considered as an alternative to the action X . $F_2(X)$ is information which relates to the doing of X by R . $F_3(X)$ is an implication, that there is some cause Z such that if Z then not-(R CAUSE X). Obviously, before the *REQUEST* body can be understood it is necessary to define the concepts implicit in it such as PROMISE, THANK, and EXPLAIN (see Appendix B) and their associated SAP's. Nevertheless, some appreciation of what the body means can be gained by an informal description.

What the SAP for *REQUEST* expresses can be stated fairly simply. In a situation in which person A WANTS person R to do something and believes that R CAN do it and that R has a reason to do it but won't in the absence of a REQUEST, A can REQUEST that the action be performed. Before we judge A's utterance as an instance of REQUEST we need to establish these preconditions and as well believe that A intends to do the asking. Regardless of the actual words uttered, it is not a REQUESTING unless this precondition of intention is established. Once we know that the situation defined by the preconditions does in fact hold then we can reasonably expect to find that A's REQUESTING exists as part of a larger pattern involving several actions. This larger pattern is given by the *REQUEST* body. It states that following the REQUESTING, R may CAUSE the action or may REFUSE. Following the REFUSAL or simply inaction of R, A may REQUEST an explanation, SUGGEST alternatives, and so on.

Other actions may also evoke *REQUEST*. For example, upon hearing, "No, I won't do that!", an observer might deduce that a REFUSE has occurred. This recognition provides a valid, though, perhaps, less likely, pointer to the *REQUEST* SAP.

In the next section we will consider a simple story and the use of SAP's, social actions and the theory of personal causation in connecting and understanding the actions.

V. The Use of Belief Systems Knowledge for Language Understanding

Consider the following story -

It was a hot July afternoon and Joey wanted to cool off. He asked his older brother Nick to buy him an ice cream cone. Nick said no. Joey then demanded an explanation. Nick said that he didn't have any money. Joey said, OK, he understood.

Part of what it means to understand this story is the ability to answer the following kinds of questions:

- Why did Joey ask Nick to buy him an ice cream cone?
- Why did Nick say he didn't have any money?
- Did Joey accept Nick's explanation for his refusal to buy the ice cream cone?

In general, these questions are not difficult, but they do require use of knowledge about social actions, e.g. REQUEST, REFUSE,

EXPLAIN, and ACCEPT (see Appendix B), patterns of behavior, e.g. *REQUEST* and related SAP's, plans, purposes, perceptions of beliefs, and other elements of belief systems knowledge. In order to illustrate how this knowledge is used in understanding language we will consider here a probable sequence of actions which would be taken by an intelligent system in understanding the story given above. We will call this system, the "Observer".

Upon hearing (or reading, seeing, etc.) the first sentence, "It was a hot July afternoon and Joey wanted to cool off," the Observer establishes a set for the story. World knowledge is brought in to make (possibly wrong) inferences such as:

- The sun is shining.
- Joey is wearing summer style clothes.
- The story is set in the Northern Hemisphere.
- Joey is perspiring.
- The temperature is over 90 degrees (F).

Using knowledge about typical behavior the Observer might also infer:

- Joey will soon do something to cool off. He may go swimming, look for an air conditioned building, sit on a block of ice, or get something cold to eat or drink.

The setting can be considered to be a collection of facts of the kind given above, the most interesting of which are the ones

which refer to beliefs and intentions. In this story we have at least one such fact, namely that Joey wants to cool off. In our notation we could say-

Joey WANT (A CAUSE X)

where X is an action which has as an eventual outcome-

Joey is-cool

and A is some person, possibly Joey. The presence of this fact implies a plan in formation. While the Observer cannot at this time know the fine structure of Joey's plan he does know certain gross characteristics. He can expect Joey to perform an action which either has the outcome, "Joey is-cool", or enables either Joey or someone else to perform some other action which ultimately has that outcome. ("Ultimately" here means the recursive application of an enablement procedure.) Since no actions are explicitly mentioned in the first sentence the Observer does not need to use social action definitions to interpret the sentence, nor are any SAP's applicable.

With the setting established, the Observer is already prepared for certain actions. Upon hearing the second sentence, "He asked his older brother Nick to buy him an ice cream cone", the Observer attempts an interpretation in terms of the structure he has constructed.

Starting with the verb (i.e. the action) he has a pointer, via "instance" links, to various social action concepts. In this case "ask" is a typical word used to express COMMAND, DEMAND, REQUEST, and perhaps other speech acts. Checking the preconditions for these concepts he finds that COMMAND is unlikely since there is no reason to think that Joey DOMINATES (has authority over) Nick. Similarly DEMAND is improbable because no OUGHT conditions have been established. REQUEST appears likely since it is reasonable to believe that each of its preconditions are satisfied.

The argument for the interpretation of Joey's action as a REQUEST might go as follows: The Observer knows from general world knowledge that-

Nick BUY an ice cream cone for Joey

ENABLES

Joey EAT ice cream cone

HAS-AS-OUTCOME

Joey is-cool

and is likely to construct such a sequence from his knowledge of Joey's wants. Furthermore the sequence provides an instantiation of the previously vague plan ascribed to Joey. Thus, assuming a

SENTIMENT and UNIT relationship between Joey and Nick it is reasonable to infer that

Joey BELIEVES (Nick TRY
(Nick BUY ice cream cone for Joey))

Also, since there is no reason to think that Nick is already planning to buy the ice cream cone, the Observer can infer that

Joey BELIEVES (Nick MUST-CHOOSE
not-(Nick BUY ice cream cone for Joey))

He already knows that

Joey WANTS (A CAUSE X)

where X leads to Joey-is-cool. From general world knowledge he might infer that since Nick is older Joey knows that he is able to buy the ice cream cone, i.e.

Joey BELIEVES Nick CAN
(Nick BUY ice cream cone for Joey)

Since the ice cream cone has not been purchased he also knows that the temporal precondition for REQUEST is satisfied, i.e.

time-request < time-action

Thus it is reasonable to infer that Joey's utterance constitutes a REQUEST.

At this point the Observer can invoke the SAP for REQUEST which we have called, *REQUEST* (see Figure 4). The header for *REQUEST* checks to see that a REQUEST has occurred and then binds its formal parameters to the case fillers of the REQUEST. In this example, the agent is Joey, the recipient is Nick, and the action is (Nick BUY ice cream cone for Joey).

REQUEST is a very general SAP, unlike, for instance, a waiter/customer pattern in a restaurant. It suggests several rather diverse actions which can reasonably be expected to follow the REQUEST. The Observer is thus predisposed though not committed to perceive successive actions to be consistent with the *REQUEST* body. Upon hearing the next sentence, "Nick said no", he tries first to interpret the action as a PROMISE, REQUEST, or other action suggested by the SAP body. In this case the negative statement, plus the fact that no other SAP is active at this point, suggest that Nick is REFUSING to fulfill Joey's REQUEST.

In order to accept Nick's action as a REFUSE, the Observer is committed either to believing that

Nick BELIEVES not-(Nick CAN
(Nick BUY ice cream cone for Joey))

or that

Nick MUST-CHOOSE

not-(Nick BUY ice cream cone for Joey)

If he cannot accept the former assumption then he must formulate a plan for Nick which is hindered by the REQUESTED action. For example, he might assume that Nick has a limited amount of money which he wants to spend on something else. As in the previous case of Joey's plan to cool off this vague plan becomes an active agent which influences the perception of succeeding sentences.

The next sentence, "Joey then demanded an explanation", fits into the SAP for *REQUEST* as a DEMAND. In fact, the interpretation of Nick's "no" as a REFUSE is reinforced by the interpretation of Joey's sentence as a DEMAND for an explanation. The combination of these interpretations suggests strongly that the next action may be an EXPLAIN by Nick. In fact, "Nick said that he didn't have any money", is a quite plausible explanation. It is an assertion that

Nick CAN (Nick BUY ice cream cone for Joey)

is not true. The final sentence, "Joey said, OK, he understood", can then be interpreted as an ACCEPT for Nick's explanation.

The description above leaves out many important details of the Observer's processing. Each action needs to be analyzed in terms of active SAP's but also in terms of formulated but unfulfilled plans, instance pointers and other possible social actions. Furthermore, the usual syntactic and semantic processing is only supplemented, not replaced by, the belief systems processing. What I have tried to show here is that intentionality and related concepts, as expressed in social action definitions and SAP's, are critical for understanding and that feasible mechanisms for their application do exist.

It should be emphasized that only in restricted cases can we do without the full machinery discussed here. We might for example explain the processing of many stories in terms of simple pattern matching rules without resort to beliefs about intentions. However, such rules fail in generality. Consider, for example, a rule which says:

- Following a question by A concerning a future action of R, a "no" from R means that R REFUSES to act.

That such a rule is insufficient can be seen clearly in the following examples:

(A)Joey: Will you buy me an ice cream cone?
Nick: No (REFUSE)

Joey: Will you forget my birthday?
Nick: No (PROMISE)

(B)Joey: Can you buy me an ice cream cone?

Nick: No (REFUSE)

Joey: Can you buy a new car?

Nick: No (INFORM)

Nick's "no" can have quite different meanings depending upon the Observer's analysis of the goals and beliefs of Nick and Joey. In the sentences in group (A) above it is an inference about Joey's WANTS and Nick's BELIEFS about Joey's WANTS which transforms the REFUSE into a PROMISE. In (B) it is an inference about the BELIEFS of Nick and Joey concerning Nick's financial capabilities which makes one "no" a likely REFUSE and the other a likely INFORM. Only by a model of beliefs and intentions can an Observer account for behavior in these social situations.

REFERENCES

Abelson, Robert

1973 "The Structure of Belief Systems", in Kenneth Colby and Roger Schank (Ed.) Computer Simulation of Thought and Language (San Francisco: W.H. Freeman & Co.).

1974 "Frames for Understanding Social Actions", in Daniel G. Bobrow and Allan M. Collins (Eds.) Representation and Understanding: Studies in Cognitive Science (New York: Academic Press).

Abelson, Robert, and J.D. Carroll

1965 "Computer Simulation of Individual Belief Systems", American Behavioral Scientist 8, 24-30.

Abelson, Robert, and Carol Reich

1969 "Implicational Molecules: A Method for Extracting Meaning from Input Sentences", Proc. First International Joint Conference on Artificial Intelligence, Wash. D.C.

Anscombe, G.E.H.

1958 "On Brute Facts", Analysis 18.

Austin, J. L.

1965 How to Do Things with Words (New York: Oxford University Press).

Brown, Geoffrey

1974 "The BELIEVER System", Computer Science Department, Rutgers, NIH Report CBM-TR-33.

Brown, Geoffrey, Bertram Bruce, and Michael Trigoboff

1974 "The CHRONOS Natural Language Understanding System", Computer Science Department, Rutgers, NIH Report CBM-TR-30.

Bruce, Bertram

- 1974 "Case Systems for Natural Language", Computer Science Department, Rutgers, NIH Report CBM-TR-31.

Bruce, Bertram and C.F. Schmidt

- 1974 "Episode Understanding and Belief Guided Parsing", Computer Science Department, Rutgers, NIH Report, CBM-TR-32.

Carbonell, Jaime and Allan Collins

- 1973 "Natural Semantics in Artificial Intelligence", Proc. Third International Joint Conference on Artificial Intelligence, Stanford.

Charniak, Eugene

- 1972 "Towards a Model of Children's Story Comprehension", MIT AI Lab Report AI-TR-256.

Colby, Kenneth Mark and David Canfield Smith

- 1969 "Dialogues Between Humans and an Artificial Belief System", Proc. First International Joint Conference on Artificial Intelligence, Wash. D.C.

Colby, Kenneth Mark, Lawrence Tesler, and Horace Enea

- 1969 "Experiments with a Search Algorithm for the Data Base of a Human Belief Structure", Proc. First International Joint Conference on Artificial Intelligence, Wash. D.C.

Colby, Kenneth, S. Weber and F.D. Hilf

- 1971 "Artificial Paranoia," Artificial Intelligence 2, 1-25.

Fillmore, Charles J.

- 1971 "Verbs of Judging: an Exercise in Semantic Description", in Charles J. Fillmore and D. Terence Langendoen (Ed.) Studies in Linguistic Semantics (New York: Holt, Rinehart and Winston, Inc.).

Garner, Richard

- 1971 "'Presupposition' in Philosophy and Linguistics", in Charles J. Fillmore and D. Terence Langendoen (Ed.) Studies in Linguistic Semantics (New York: Holt, Rinehart and Winston, Inc.).

Gordon, David, and George Lakoff

- 1971 "Conversational Postulaces", Proc. Seventh Regional Meeting, Chicago Linguistic Society.

Grice, H.P.

- 1957 "Meaning", Philosophical Review 66, 377-388.

Hass, Wilbur

- 1971 "Truth-Functional and Communicational Bases for Prescriptive Discourse", Proc. Seventh Regional Meeting, Chicago Linguistic Society.

Heider, Fritz

- 1958 The Psychology of Interpersonal Relations (new York: Wiley).

Hutchinson, Larry G.

- 1971 "Presupposition and Belief-Inferences", Proc. Seventh regional meeting, Chicago Linguistic Society.

Lakoff, George

- 1971 "Presupposition and Relative Well-Formedness", in D.D. Steinberg and L.A. Jakobovits (Eds.) Semantics: An Interdisciplinary Reader (Cambridge: Cambridge University Press).

McCarthy, John and P.J. Hayes

- 1969 "Some Philosophical Problems from the Standpoint of Artificial Intelligence" in B. Meltzer and D. Michie (Eds.) Machine Intelligence 4 (New York: American Elsevier).

McDermott, Drew

- 1974 "Assimilation of New Information by a Natural Language Understanding System", MIT AI Lab Report AI-TR-291.

Minsky, Marvin

- 1968 "Matter, Mind, and Models", in Marvin Minsky (Ed.) Semantic Information Processing (Cambridge, Mass.: The MIT Press).

- 1974 "A Framework for Representing Knowledge", to appear in The Psychology of Computer Vision (New York: McGraw-Hill).

Norman, Donald

- 1972 "Memory, Knowledge, and the Answering of Questions", Center for Human Information Processing, U. of California, San Diego, Chip 25.

Rumelhart, David E.

- 1974 "Notes on a Schema for Stories", in Daniel G. Bobrow and Allan M. Collins (Eds.) Representation and Understanding: Studies in Cognitive Science (New York: Academic Press).

Schank, Roger

- 1972 "Conceptual Dependency: A Theory of Natural Language Understanding", Cognitive Psychology 3, 552-631.

- 1973 "Causality and Reasoning", Istituto per gli studi Semantici e Cognitivi, Castagnola, Switzerland, TR-1.

Schank, Roger, Neil Goldman, Charles J. Rieger, and Christopher K. Riesbeck

- 1972 "Primitive Concepts Underlying Verbs of Thought", Computer Science Department, Stanford, STAN-CS-265-72.

Schmidt, C.F.

- 1974 "Recognizing Plans and Purposes", Computer Science Department, Rutgers, NIH Report, CBM-TR-34.

Schmidt, C.F. and John D'Addamio

- 1973 "A Model of the Common-Sense Theory of Intention and Personal Causation", Proc. Third International Joint Conference on Artificial Intelligence, Stanford.

Sedlak, Andrea

- 1974 "An Investigation of the Development of the Child's Understanding and Evaluation of the Actions of Others", Computer Science Department, Rutgers, NIH Report, CBM-TR-28.

Searle, J.R.

- 1969 Speech Acts: An Essay in the Philosophy of Language (Cambridge: Cambridge University Press).

Stravson, P.F.

1964 "Intention and Convention in Speech Acts", Philosophical Review 73, 439-460.

Wittgenstein, Ludwig

1958 Philosophical Investigations (New York: The Macmillan Co.) (Tr. by G.E.M. Anscombe).

ACKNOWLEDGEMENTS

I owe many of the ideas discussed here to Chuck Schmidt and Geoff Brown. In fact it is Schmidt's theory of personal causation which provides the foundation for this paper. Geoff Brown is implementing some of these ideas in a computer program for understanding stories in natural language. (The program is called BELIEVER). I would also like to thank Andee Rubin and Eleanor Warnock for suggestions on the paper itself. Typing and preparation of this report was done by Linda Amsden.

This research was also supported by the Rutgers Special Research Resource on Computers in Biomedicine.

APPENDIX A - Notation Conventions

The language used here is a form of predicate calculus (see Bruce and Schmidt (1974)). In order to distinguish social actions from the English words which are commonly used to describe them, the predicates in our language are written in capital letters. Thus, the concept of an "informing" is represented by the symbol "INFORM" rather than the English word "inform". In a similar way concepts such as "being able to", "being motivated to", "being obligated to", and "believing" are represented by the symbols "CAN", "TRY", "OUGHT", and "BELIEVE".

For the sake of readability several deviations from standard predicate calculus notation will be used. Rather than expressing atomic formulas in the form:

<predicate> (<argument 1> ... <argument n>)

we will use the more English-like form:

<argument 1> <predicate> <argument 2> ... <argument n>

Thus INFORM (A, B, S, t) (meaning "A informs B that S at time t") will be written:

A INFORM B S t

In addition various function symbols such as "that", "to", "-s", and "does" will be used whenever they make the formulas more readable. For example:

A INFORMS B that S at t

Where no confusion can result the explicit reference to times will be omitted. Thus,

A INFORMS B that S .

APPENDIX B - Social Actions

ACCEPT

(Frances accepted Ralph's explanation for being late)

Case Structure

agent: A
recipient: R
cause: X
action: Y
time-accept: t
time-action: t'

Preconditions

- P1. A MUST-CHOOSE (A ACCEPT R X)
- P2. A BELIEVES (R EXPLAIN A X Y)
- P3. A BELIEVES (if X then not-(R CAUSE Y))
- P4. $t' < t$

Outcome Conditions

- O1. CONSEQUENCES (PC + OC)

Instances accept, understandSummary Characterizations accept

FIGURE 5 The social action, ACCEPT. Precondition P2 states that in order for A's action to be an ACCEPT A must BELIEVE that R has EXPLAINED why R did not CAUSE a previous action, Y. The explanation is some proposition, X. P3 states that A BELIEVES the explanation, i.e. that X implies either not-(R CAN Y) or not-(R MUST-CHOOSE Y).

ACCUSE

(George accused Bill of dealing from the bottom of the deck)

Case Structure

agent: A
recipient: R
action: X
time-accept: t
time-action: t'

Preconditions

- P1. A MUST-CHOOSE (A ACCUSE R X)
- P2. A BELIEVES (R CAUSE X)
- P3. A BELIEVES (R SUFFICIENT-UGHT not-X)
- P4. $t' < t$

Outcome Conditions

- O1. CONSEQUENCES (PC + OC)

Instances accuse, blameSummary Characterizations accuse, blame

FIGURE 6 The social action, ACCUSE. An ACCUSE occurs when A BELIEVES that R CAUSED some action, X, (P2) and that R had a sufficient obligation not to do X. It can be followed by a form of EXPLAIN which presents reasons why R CAUSED X, by an action such as APOLOGIZE, another ACCUSE, a DENY or a REINTERPRET.

ASKABOUT

(Alfred asked Judy about the upcoming schedule)

Case Structure

agent: A
recipient: R
object: X
time-askabout: t

Preconditions

- P1. A MUST-CHOOSE (A ASKABOUT X to R)
- P2. not-(A KNOWS(is-aware-of) X)
- P3. A BELIEVES (R CHOOSE (R INFORM A X t'))
- P4. $t < t'$
- P5. A BELIEVES (R KNOWS(is-aware-of) X)

Outcome Conditions

- 01. CONSEQUENCES (PC + OC)

Instances who, what, where, etc.; tell me..., is..., did...

Summary Characterizations ask

FIGURE 7 The social action, ASKABOUT. Precondition P2 says that the agent of ASKABOUT does not know the answer to his question. If P2 is false then R may respond, "I thought you knew!". P2 being false suggests that the apparent ASKABOUT is really a TEST, a form of LIE, or ritual talk. However it could also be that the agent is trying to clarify or strengthen his knowledge.

If P5 is false then instead of A ASKABOUT X to R we have A ASKABOUT (R KNOWS(is-aware-of) X). If P3 is false then the action could be an attempt to EMBARRASS or BADGER.

COMMAND

(Alexander commanded his troops to shave their beards)

Case Structure

agent: A
 recipient: R
 action: X
 time-command: t
 time-action: t'

Preconditions

- P1. A MUST-CHOOSE (A COMMAND R to X)
- P2. A WANTS (R CAUSE X)
- P3. A BELIEVES (R CAN X)
- P4. A BELIEVES (R CHOOSE not-X)
- P5. A BELIEVES (R SUFFICIENT-UGHT X)
- P6. A BELIEVES (if (R BELIEVES (A COMMAND R to X))
 then (R OUGHT X))
- P7. A DOMINATES R
- P8. $t < t'$

Outcome Conditions

- O1. R OUGHT X
- O2. CONSEQUENCES (PC + OC)

Instances order, command, tell

Summary Characterizations order, command

FIGURE 3 The social action, COMMAND. COMMAND is a complex form of "asking". In addition to the preconditions that A WANTS (R CAUSE X) and A BELIEVES (R CAN X) there are conditions on the relationship of A to R and on A's beliefs about R's intentions.

A must BELIEVE that R has reason(s) not to do X but that there is a SUFFICIENT-UGHT for R to do X. Furthermore A BELIEVES that if R recognizes the COMMAND that R OUGHT X. Finally, (in this situation) A DOMINATES R.

COMMAND has a special outcome condition (O1) which says that the observer (anyone who recognizes A's action as a COMMAND) BELIEVES that R OUGHT X.

CONVINCE

(Peter convinced me to buy the blueberries)

Case Structure

agent: A
recipient: R
proposition: X
time-convince: t

Preconditions

- P1. A MUST-CHOOSE (A CONVINCE R that X)
- P2. A BELIEVES X
- P3. A BELIEVES not-(R BELIEVES X)
- P4. not-(R BELIEVES X)

Outcome Conditions

- O1. R BELIEVES S
- O2. CONSEQUENCES (PC + OC)

Instances tell, inform, say

Summary Characterizations convince

FIGURE 2 The social action, CONVINCE. CONVINCE differs from INFORM in that it implies a stronger commitment about the beliefs of the recipient. Precondition P4 states that R does not BELIEVE the proposition X, a constraint which is not implied by a simple INFORMING. Also CONVINCE makes a stronger statement about the effects of the action. Whereas INFORM says merely that the recipient is aware of the proposition, CONVINCE says that the recipient must BELIEVE the proposition.

There is another sense of "convince" which does not express intentionality; instead it merely gives the outcome of a discussion. Thus we can talk of "convincing" in the intentional sense (like PROMISE) or the descriptive sense. Many other words, such as "ask" serve similar dual roles. In this paper we are concerned with the intentional uses only.

DEMAND

(She demanded that he shut up)

Case Structure

agent: A
recipient: R
action: X
time-demand: t
time-action: t'

Preconditions

- P1. A MUST-CHOOSE (A DEMAND R to X)
- P2. A WANTS (R CAUSE X)
- P3. A BELIEVES (R CAN X)
- P4. A BELIEVES (R BELIEVES (R MUST-CHOOSE not-X))
- P5. A BELIEVES (R SUFFICIENT-UGHT X)
- P6. $t < t'$

Outcome Conditions

- O1. if not-(R CAUSE X)
then R OUGHT (R EXPLAIN A Z X)
- O2. CONSEQUENCES (PC + OC)

Instances demand, insistSummary Characterizations demand, insist

FIGURE 10 The social action, DEMAND. DEMAND is a complex form of "asking". In addition to the preconditions that A WANTS (R CAUSE X) and A BELIEVES (R CAN X) there are preconditions that A BELIEVES (R BELIEVES (R MUST-CHOOSE not-X) and A BELIEVES that R intends not to do X but that there is a strong OUGHT for R to do X. Because A's belief in the SUFFICIENT-UGHT there is a local normative rule invoked by DEMAND. This rule, pointed to by O1, states that if R does not do X then R OUGHT to give a reason why (EXPLAIN).

EXPLAIN

(Jerry explained to Jane that the cake burned because the thermostat was broken)

Case Structure

agent: A
 recipient: R
 cause: X
 action: Y
 time-explain: t
 time-cause: t'
 time-action: t''

Preconditions

- P1. A MUST-CHOOSE (A EXPLAIN to R
 that (if X then not-(A CAUSE Y)))
- P2. not-(A CAUSE Y)
- P3. A BELIEVES (R BELIEVES not-(A CAUSE Y))
- P4. A BELIEVES (R WANTS (A CAUSE Y))
- P5. A BELIEVES (if X then not-(A CAUSE Y))
- P6. A BELIEVES (R BELIEVES
 (if X then not-(A CAUSE Y)))
- P7. A BELIEVES not-(A OUGHT Y)
- P8. t' < t''
- P9. t'' < t

Outcome Conditions

- O1. CONSEQUENCES (PC + OC)

Instances explainSummary Characterizations explain, make an excuse

FIGURE 11 The social action, EXPLAIN. A variant of this social action is one in which the agent is EXPLAINING why an unwanted act was done (rather than a wanted act which was not done). In either case the cause, X, must be a sufficient reason for A's behavior. In the case of an act not done X should say why A CAN X is not true or why A MUST-CHOOSE X is not true. If X is focused on MUST-CHOOSE then it can be a statement that A has a sufficient reason not to do the act, either A SUFFICIENT-TRY not-X or A SUFFICIENT-UGHT not-X. X can also state that there is simply no sufficient reason to do the act. The exact form which X takes in a specific situation depends upon such things as the previous social actions. For example, if A has failed to carry out a valid COMMAND then the only appropriate causes for the lack of action would be not-(A CAN X) or A SUFFICIENT-UGHT not-X.

INFORM

(Linda informed Richard that the tire was flat)

Case Structure

agent: A
recipient: R
proposition: X
time-inform: t

Preconditions

- P1. A MUST-CHOOSE (A INFORM R that X)
- P2. A BELIEVES X
- P3. A BELIEVES not-(R BELIEVES X)

Outcome Conditions

- O1. R KNOWS(is-aware-of) S
- O2. if not-(R BELIEVES S)
then R OUGHT (R INFORM A not-(R BELIEVE S))
- O3. CONSEQUENCES (PC + OC)

Instances tell, saySummary Characterizations inform

FIGURE 12 The social action, INFORM. If P2 is false then the apparent INFORM is probably a form of ritual talk or a LIE. If P3 is false then it is probably a REQUEST or COMMAND. If X is an action such that A causes a negative outcome for R then it may be a THREAT; if the outcome is positive then it may be a PROMISE. If X contains a "need" or "want" verb then the action may be a REQUEST or COMMAND for the complement of the need-want verb. If X asserts that A cannot do an action Y then it may be a REQUEST for R to do Y. If X asserts that A doesn't have some object Z then it may be a REQUEST for HELP from R in obtaining Z.

There is probably another related social action, which might be called "TELL". TELL has the same outcome conditions and instances as INFORM. It differs only in not requiring P3, i.e. there is no commitment made about R's beliefs. While the notion of a simple TELLING is useful in some circumstances, it should not be used to hide more motivated speech acts such as REQUESTING which often are expressed as simple factual statements.

LIE

(Sam lied to Karen about the money)

Case Structure

agent: A
recipient: R
proposition: X
time-lie: t

Preconditions

- P1. A MUST-CHOOSE (A LIE X to R)
- P2. not-(A BELIEVE X)
- P3. A BELIEVES not-(R BELIEVES not-X)

Outcome Conditions

- O1. CONSEQUENCES (PC + OC)

Instances tell, say, inform

Summary Characterizations lie

FIGURE 13 The social action, LIE. The preconditions are similar to those for INFORM, except, of course, that a "not" is inserted. P3 says that the agent of LIE must BELIEVE that the recipient does not BELIEVE the opposite of the proposition, i.e., A BELIEVES there is some chance of success with the LIE. If R BELIEVES that an INFORMING has occurred then the consequences of that social action will hold within R's belief system. If P1 is false then A can say later that he was mistaken (P2) and did not intend to LIE. If R BELIEVES not-X then R may say "You are wrong" or "Did I hear you correctly?"

PROMISE

(Catherine promised to give Bill a kiss)

Case Structure

agent: A
recipient: R
action: X
time-promise: t
time-action: t'

Preconditions

- P1. A MUST-CHOOSE (A PROMISE R X)
- P2. A CAN X
- P3. R WANTS (A CAUSE X)
- P4. if not-(A PROMISE R X)
then not-(A MUST-CHOOSE X)
- P5. A TRY X
- P6. $t < t'$

Outcome Conditions

- 01. A SUFFICIENT-UGHT X
- 02. CONSEQUENCES (PC +OC)

Instances promise, offerSummary Characterizations promise

FIGURE 14 The social action, PROMISE. PROMISE is a rather complicated speech act and its definition makes some fine distinctions. Basically the preconditions say that A is capable of doing X and intends to do it; that he would not necessarily do it were it not for the PROMISE and that R WANTS A to do X. The important outcome condition is that A now has a sufficiently strong obligation to carry out the action PROMISED. For a further discussion of PROMISE with a slightly different approach, see Searle (1969).

REFUSE

(Art refused to divulge the location of the treasure)

Case Structure

agent: A
recipient: R
action: X
time-refuse: t

Preconditions

- P1. A MUST-CHOOSE (A REFUSE R X)
- P2. A BELIEVES (R ASK A to X)
- P3. A MUST-CHOOSE not-X
- P4. not-(A CAUSE X)

Outcome Conditions

- O1. CONSEQUENCES (PC + JC)
- O2. R BELIEVES not-(A CAUSE X)

Instances refuse, I won't ...

Summary Characterizations refuse

FIGURE 15 The social action, REFUSE. REFUSE is an action which typically follows an ASKING...TO concept such as ASK, DEMAND, or COMMAND. It occurs when A BELIEVES that the ASK has happened but MUST-CHOOSE not to act. One outcome of REFUSE is that R BELIEVES that A will not do X.

REQUEST

(Mac requested that Ella sharpen the knife)

Case Structure

agent: A
recipient: R
action: X
time-request: t
time-action: t'

Preconditions

- P1. A MUST-CHOOSE (A REQUEST R to X)
- P2. A WANTS (R CAUSE X)
- P3. A BELIEVES (R CAN X)
- P4. A BELIEVES (R TRY X))
- P5. A BELIEVES (R MUST-CHOOSE not-X)
- P6. $t < t'$

Outcome Conditions

- 01. if not-(R CAUSE X)
then R OUGHT (R EXPLAIN A Z X)
- 02. CONSEQUENCES (PC + OC)

Instances will you ..., could you ...

Summary Characterizations request, ask

FIGURE 16 The social action, REQUEST. REQUEST is a type of "asking" which implies beliefs of A about R's intentions. A BELIEVES that R has a reason to do X but that MUST-CHOOSE is not true. The main purpose of the REQUEST is to change R's choosevalues so that R MUST-CHOOSE X will be true.

TEST

(Stan tested Nell's knowledge of lacrosse)

Case Structure

agent: A
recipient: R
object: X
time-test: t

Preconditions

- P1. A MUST-CHOOSE (A TEST X to R)
- P2. A BELIEVES X
- P3. A BELIEVES (R CHOOSE
((R INFORM A X t') or (R INFORM A not-X t')))
- P4. $t < t'$
- P5. not-(A BELIEVES (R BELIEVES X))
- P6. not-(A BELIEVES (R BELIEVES not-X))

Outcome Conditions

- O1. CONSEQUENCES (PC + OC)

Instances who, what, where, etc.; tell me..., is..., did...

Summary Characterizations test, ask

FIGURE 17 The social action, TEST. Preconditions P5 and P6 state that the agent does not BELIEVE either that R BELIEVES or does not BELIEVE X. Thus his asking is an attempt to ascertain part of R's beliefs.

THANK

(Frank thanked Nancy for the volleyball)

Case Structure

agent: A
recipient: R
action: X
time-thank: t
time-action: t'

Preconditions

- P1. A MUST-CHOOSE (A THANK R X)
- P2. A BELIEVES (R CAUSE X)
- P3. A WANT (R CAUSE X)
- P4. $t' < t$

Outcome Conditions

- O1. CONSEQUENCES (PC + OC)
- O2. R KNOW(is-aware-of) (A WANT (R CAUSE X))

Instances thankSummary Characterizations thank

FIGURE 18 The social action, THANK. THANK occurs when A BELIEVES that R has done an action which A WANTS. One consequence is that R is aware of the fact that A wanted R to do the action.

THREATEN

(Norman threatened to leave the party if Pat wouldn't stop singing)

Case Structure

agent: A
 recipient: R
 action: X
 r-action: Z
 time-threaten: t
 time-action: t'
 time-r-action: t''

Preconditions

- P1. A MUST-CHOOSE (A THREATEN X to R)
- P2. A CAN X
- P3. A BELIEVES (R BELIEVES (A CAN X))
- P4. A BELIEVES not-(R WANT (A CAUSE X))
- P5. A BELIEVES (R CAN Z)
- P6. A BELIEVES (R CHOOSE Z)
- P7. A WANT not-(R CAUSE Z)
- P8. A BELIEVES (R BELIEVES not-(A WANT (R CAUSE Z)))
- P9. $t < t''$
- P10. $t'' < t'$

Outcome Conditions

- 01. CONSEQUENCES (PC + OC)

Instances you better not ..., I warn you

Summary Characterizations threaten

FIGURE 19 The social action, THREATEN. THREATEN is a complex social action which is related to the "asking" concepts because it involves A's seeking to influence R's action. Most of the preconditions fall into two categories, those relating to A's THREATENED action (P2-P4) and those relating to R's action (P5-P8).

WARN

(Leslie warned Fran about the troll)

Case Structure

agent: A
recipient: R
action: X
time-warn: t
time-action: t'

Preconditions

- P1. A MUST-CHOOSE (A WARN X to R)
- P2. A BELIEVES not-(R WANT (OUTCOMES X))
- P3. A BELIEVES
(some Z)((R CAN Z) and
(if (R CAUSE Z) then not-X))
- P4. A SENTIMENT R or A UNIT R
- P5. $t < t'$

Outcome Conditions

- O1. CONSEQUENCES (PC + OC)

Instances you better not, I warn youSummary Characterizations warn

FIGURE 20 The social action, WARN. WARN is an action whereby A attempts to get R to act in such a way that R avoids actions which A BELIEVES have OUTCOMES which R does not WANT. A must have some SENTIMENT (e.g. "likes") or UNIT (e.g. "family") relationship to R.